

Privacy Governance-Driven Design of AI-Powered Elderly Safety Monitoring for Cambodia

Monireach Tang*, Seingheng Hul**, May Thu***

Abstract Cloud-based elderly monitoring systems create surveillance risks with limited data protection enforcement. This study validates technical feasibility and cost-effectiveness of edge-based pose estimation for elderly safety monitoring in resource-constrained contexts, demonstrating how privacy principles can drive architectural decisions from inception. The proposed architecture comprises four 850nm NIR cameras and an NVIDIA Jetson Orin Nano edge platform. Privacy protection emerges from system constraints: immediate conversion of video frames to skeletal coordinates, pose-only storage, and permanent video deletion. The integrated YOLOv8n and MediaPipe pipeline achieves 100% on-device processing. Validation on 20 commercial CCTV videos demonstrates 91.3% keypoint detection at 20.53 FPS, confirming 24/7 monitoring capability without facial recognition technology. Edge architecture reduces 3-year total cost by 61% compared to cloud alternatives, expanding market accessibility to 168,000–252,000 elderly individuals in middle-income Cambodian urban households. This work demonstrates that governance principles can drive technical architecture in healthcare AI for Cambodia.

Keywords Privacy governance, edge computing, elderly safety monitoring, healthcare AI, privacy-by-design

Submitted, December 30, 2025; Accepted, March 06, 2026

* Postgraduate Student, School of Graduate Studies, Cambodia University of Technology and Science, Phnom Penh, Cambodia; monireach.tang@camtech.edu.kh

** Under Secretary of State, Ministry of Industry, Science, Technology & Innovation, Phnom Penh, Cambodia; hul.seingheng@misti.gov.kh

*** Dean, Faculty of Engineering, Cambodia University of Technology and Science, Phnom Penh, Cambodia; may.thu@camtech.edu.kh



This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.

I. Introduction

1. Problem Context

Since the 2000s, global demographic transitions have constituted a pressing challenge for healthcare systems. Southeast Asia's elderly population is projected to increase from 12.2% in 2024 to 22.9% by 2050 (WHO 2025), representing a fundamental restructuring of regional age distributions. Cambodia's elderly population (60+ years) will reach approximately 2.1 million by 2030, comprising 11% of the projected 19 million total population (United Nations, Department of Economic and Social Affairs, Population Division 2015), with 8-12% residing in middle-income urban households (4th-5th income quintiles).

Falls represent the leading cause of injury-related deaths among elderly populations—684,000 annual fatalities globally, with 60% concentrated in Western Pacific and Southeast Asia regions (WHO 2021). The challenge that middle-income families face is providing continuous supervision for elderly relatives while managing employment obligations. Strong family caregiving norms rooted in filial piety values create a preference for home-based elderly care over institutional facilities across Southeast Asia (Romli et al. 2017), which means that technology solutions must fit within family homes rather than institutional deployment contexts. This cultural context shapes both the problem space and solution constraints in ways that Western healthcare AI research does not always adequately consider.

2. Technology Governance Challenges

Existing elderly safety monitoring technologies present critical governance challenges across two dimensions: privacy concerns and accessibility constraints.

Privacy governance challenges. Cloud-based AI systems transmit sensitive health data—video footage of individuals in private home environments, behavioral patterns, movement trajectories—to third-party servers for processing and storage. This design choice creates facial recognition risks, behavioral profiling capabilities, and re-identification vulnerabilities that affect elderly populations who may not fully understand surveillance implications (Burns 2024; Birkstedt et al. 2023). Commercial fall detection camera systems exemplify these concerns, requiring continuous cloud connectivity for fall detection processing (Upadhyay 2025), which creates the privacy exposures that governance frameworks caution against.

Data protection regulations exist, but enforcement mechanisms remain limited in Southeast Asian contexts (Burns 2024). Privacy protection cannot rely solely on regulatory compliance but must emerge from system architecture itself. This brings us to Privacy by Design principles (Cavoukian, Taylor, and Abrams 2010), which emphasize embedding privacy into architecture from inception rather than retrofitting protections after deployment.

Accessibility governance challenges. For long-term affordability, cloud-based systems require ongoing subscription fees that prove more consequential than initial hardware costs. The Kami Fall Detect Camera charges approximately \$45 monthly (\$540 annually) in addition to initial hardware costs (Upadhyay 2025). For middle-income Cambodian households earning \$870-\$1,622 monthly (National Institute of Statistics (Cambodia) 2021), this subscription represents 2.8-5.2% of monthly income, transforming a technology purchase into an ongoing financial obligation that competes with essential household expenses.

This economic barrier undermines digital inclusion objectives. Sylla, Ismaila, and Diallo (2025) document persistent economic and infrastructure barriers to technology adoption in low- and middle-income countries, while Richardson et al. (2022) propose Digital Determinants of Health frameworks that acknowledge how cost structures systematically exclude populations from healthcare technology benefits.

Governance fragmentation. AI governance frameworks remain nascent and fragmented (Birkstedt et al. 2023; Burns 2024; Sylla, Ismaila, and Diallo 2025). The field has produced extensive literature on ethical principles—transparency, accountability, fairness, explainability—while providing limited practical guidance for operationalizing these principles in resource-constrained healthcare contexts. Almeida, Santos, and Farias (2020) synthesize 21 AI governance models into a unified framework, though the synthesis of high-level principles does not directly translate to architectural specifications. This gap between principle and practice motivates the governance-driven design approach demonstrated in this paper.

3. Research Gap

Existing elderly monitoring solutions prioritize either privacy protection, affordability, or technical performance—rarely addressing all three dimensions as integrated design objectives (Chang-Yueh Wang and Fang-Suey Lin 2024; Jalal, Kamal, and Kim 2014; Vaiyapuri et al. 2021). This fragmentation reflects that privacy, cost, and performance optimization respond to different design incentives that research and commercial development don't naturally integrate.

Cavoukian, Taylor, and Abrams (2010) Privacy by Design framework articulates philosophical principles but lacks empirical validation in healthcare AI contexts showing how these principles translate to specific architectural choices. Burns (2024) identifies edge computing potential for privacy governance through data locality but does not quantify cost-effectiveness implications for developing countries, where cloud infrastructure costs create accessibility barriers. Almeida, Santos, and Farias (2020) examine healthcare AI ethics but focus on clinical decision-making systems rather than continuous monitoring applications for vulnerable populations in residential settings.

Edge computing's potential for enabling privacy governance in healthcare AI remains underexplored, particularly regarding whether privacy-first architectural design yields economic co-benefits that expand market accessibility. The question is not whether edge computing can enhance privacy (data that never leaves local devices cannot be compromised by cloud breaches), but whether edge architecture simultaneously reduces costs in ways that matter for middle-income market segments. Existing literature treats privacy and affordability as separate optimization dimensions rather than examining whether architectural choices that enforce privacy might also eliminate cloud infrastructure expenses.

Regional studies on elderly care technology adoption in Southeast Asia focus primarily on high-income urban populations (Romli et al. 2017), with limited attention to middle-income market segments that represent larger potential user bases but face affordability constraints.

4. Research Objectives

This study addresses the identified gaps through three objectives:

- **Demonstrate governance-driven architectural design:** Show how privacy governance principles can drive technical architecture decisions from inception rather than being retrofitted post-deployment, using elderly safety monitoring for Cambodia as a case study.
- **Validate technical feasibility of privacy-preserving edge architecture:** Empirically test whether affordable 850nm near-infrared (NIR) security cameras maintain adequate pose estimation performance for continuous monitoring applications. This addresses the highest-uncertainty technical assumption: that pose estimation algorithms trained predominantly on visible-spectrum imagery will work on NIR footage from affordable CCTV cameras.
- **Quantify cost-effectiveness and accessibility implications:** Compare edge-based architecture total costs against cloud-dependent alternatives to

assess whether privacy-first design expands market accessibility for middle-income Cambodian households.

This work contributes empirical evidence that privacy governance need not trade off against affordability or performance. When integrated at the architectural design stage rather than applied as a post-hoc compliance measure, privacy principles can guide technical decisions that simultaneously enhance data protection and expand economic accessibility. The study targets middle-income urban Cambodian households (4th-5th income quintiles, representing an estimated 8-12% of the elderly population by 2030) but has broader implications for healthcare AI governance in resource-constrained contexts globally.

II. Literature Review

1. Elderly Safety Monitoring Technologies

Wearable devices represent the dominant fall detection approach, though compliance challenges limit their effectiveness. Elderly users must remember to wear devices consistently and maintain battery charging, tasks that prove particularly difficult at nighttime when fall risk increases. Only 7.1% of wearable device projects monitored the elderly in actual home environments, with most testing confined to controlled laboratory settings (Chaudhuri, Thompson, and Demiris 2014). The gap between lab performance and real-world deployment remains a persistent problem in assistive technology research.

Cloud-based camera systems emerged as an alternative monitoring modality. These systems transmit video to third-party servers for processing, solving the computational problem while creating substantial privacy concerns (Upadhyay 2025). The Kami Fall Detect Camera requires continuous cloud connectivity and ongoing subscription payments, presenting accessibility barriers. Moreover, elderly populations perceive video surveillance as intrusive (Uddin, Khaksar, and Torresen 2018).

Depth sensors (RGB-D cameras employing 3D imaging) have been proposed as privacy-preserving alternatives. These systems eliminate facial recognition capability by design rather than policy promise (Jalal, Kamal, and Kim 2014; Wang, Ellul, and Azzopardi 2020). However, sensor modality tradeoffs persist: privacy improvements come bundled with cost increases and performance constraints (Wang, Ellul, and Azzopardi 2020).

Ambient sensors—motion detectors, passive infrared (PIR) sensors, door contact sensors—face signal ambiguity challenges. Falls and Activities of Daily Living produce similar sensor patterns (Uddin, Khaksar, and Torresen 2018;

Wang, Ellul, and Azzopardi 2020), limiting incident-specific detection capability. Wang, Ellul, and Azzopardi (2020) report detection precision below 90% for individual sensor systems. Furthermore, Uddin, Khaksar, and Torresen (2018) observe a “lack of suitable outcomes to validate delivery of technological solutions for specific needs,” suggesting that systems have been built without adequately defining success metrics.

2. Privacy-Preserving AI Approaches

Federated learning trains models across distributed data without sharing raw information (Yu et al. 2022), representing genuine progress in privacy protection. However, federated learning still requires cloud infrastructure for coordination, along with associated costs and network exposure that undermine complete privacy claims.

Differential privacy adds statistical noise to data to protect individual privacy while enabling analysis (Liu et al. 2023; Williamson and Prybutok 2024). The technique is mathematically elegant but presents challenges in healthcare AI contexts. Liu et al. (2023) document that noise injection reduces model accuracy, which poses problems for safety-critical applications where missing a fall detection could prove fatal. Fine-tuning noise parameters to achieve optimal privacy-utility balance requires expertise and computational resources not readily available in resource-constrained healthcare contexts.

Edge computing processes data on local devices rather than transmitting to cloud servers, providing what Burns (2024) terms “data locality benefits.” Chang, Hsu, and Chen (2021) demonstrate that edge deployment eliminates cloud transmission and associated privacy concerns. Edge computing introduces its own constraints: local hardware limitations, update deployment challenges, and maintenance requirements.

Cavoukian, Taylor, and Abrams (2010) Privacy by Design framework articulates principles that emphasize embedding privacy into architecture from inception rather than retrofitting protections after system design. This distinction maps onto the difference between edge computing (privacy by architecture) and differential privacy (privacy by mathematical intervention). Cavoukian identifies seven foundational principles, including “Privacy Embedded into Design” and “Proactive not Reactive,” which prove difficult to implement in practice.

3. Technology Governance Frameworks

Almeida, Santos, and Farias (2020) synthesized 21 AI governance models into a unified framework emphasizing fairness, transparency, accountability, safety,

and human rights. Birkstedt et al. (2023) observe that the challenge lies in translating these ethical principles into organizational practice, noting that implementation guidance remains limited.

Data protection regulations represent another governance layer with uneven global implementation. GDPR mandates privacy safeguards including data controllership, right to erasure, and data minimization (Williamson and Prybutok 2024). However, enforcement mechanisms remain limited in Southeast Asian contexts (Burns 2024), which means that privacy protection must be embedded in system architecture rather than relying primarily on regulatory pressure.

Healthcare privacy standards add another dimension. Burns (2024) emphasizes that privacy and accountability principles prove essential for health security applications. Healthcare privacy standards remain heavily Western-centric, with limited frameworks developed for or by developing countries.

Accessibility and digital inclusion policies address a critical governance dimension. Richardson et al. (2022) propose a Digital Determinants of Health framework operating across individual, interpersonal, community, and societal levels. Hatef et al. (2024) examine digital health equity frameworks, while Sylla, Ismaila, and Diallo (2025) document persistent economic and infrastructure barriers to technology adoption in low- and middle-income countries (LMICs). When subscription-based monitoring systems require ongoing payments representing 5% or more of monthly household income, economic exclusion by design occurs.

4. Regional Context: Southeast Asian Elderly Care

Epidemiological evidence from Southeast Asia grounds incident type prioritization. Maiyapakdee et al. (2025) report that 37.7% of elderly home accidents in Thailand are falls, justifying fall detection as a priority safety concern. Li and Shi (2022) examined 272 nursing facilities in China, identifying falls as the top adverse event, with bed and chair falls explicitly documented. Romli et al. (2017) synthesize broader regional patterns, reporting fall prevalence of 7.5-10% in Thailand and Singapore, with ranges of 14-34% across Asia.

Cultural norms regarding elderly care prove salient for system design. Romli et al. (2017) document strong family caregiving norms rooted in filial piety values, with a preference for home-based care over institutional facilities across Southeast Asian contexts. Older adults in ASEAN communities often live in extended family households with adult children providing supervision, though families cannot provide 24/7 monitoring. Systems must fit within family homes rather than institutional settings.

Economic constraints represent a critical dimension of regional context. Cambodia's middle-income urban households (4th and 5th quintiles earning \$870-\$1,622 per month) face affordability thresholds that exclude many technologies developed for higher-income markets (National Institute of Statistics (Cambodia) 2021). Sylla, Ismaila, and Diallo (2025) describe persistent economic and infrastructure barriers to technology adoption in LMICs as concrete financial constraints: systems costing several months' income prove inaccessible regardless of technical sophistication. Hatef et al. (2024) document how cost barriers compound other Digital Determinants of Health, creating multilevel obstacles to digital health equity. This positions cost as a governance issue beyond mere economics: pricing structures determine which populations can access healthcare AI, making affordability decisions a form of governance that either enables or excludes populations from technological safety benefits.

III. Methodology

1. Privacy Governance Architecture

This study proposes a privacy governance-driven architecture for elderly safety monitoring and validates its technical feasibility through NIR camera compatibility testing. The architecture begins with privacy governance principles as primary constraints, deriving technical design from these requirements rather than retrofitting privacy protections post-development.

1.1 Privacy Governance Principles

Two architectural mechanisms enforce privacy by design (Cavoukian, Taylor, and Abrams 2010):

Edge-first processing. All computation occurs on a local device (NVIDIA Jetson Orin Nano 8GB for target deployment). Zero cloud transmission eliminates network exposure entirely. Data sovereignty concerns vanish when data never leaves the household.

Pose-only storage with immediate frame disposal. The system stores only 17 body keypoints in COCO format (Common Objects in Context standard: shoulders, elbows, wrists, hips, knees, ankles, and torso keypoints), excluding face landmarks entirely. MediaPipe's full pose model detects 33 landmarks, but facial keypoints are intentionally discarded during extraction. Raw video frames are never persisted to storage—not encrypted, not anonymized, simply never written to disk. Video frames undergo processing in real-time, with deletion occurring immediately after keypoint extraction. The temporal gap between frame capture and deletion measures in milliseconds, limited by processing

pipeline speed. This design choice eliminates facial recognition capability structurally: no facial data exists to enable identification, and no video retention eliminates retrospective re-identification risk, data breach exposure of video archives, and dependency on deletion policy compliance.



Source: Adapted from 180 Degree HD Dome Security Camera with Infrared Night Vision by CCTV Camera Pros (YouTube). Annotations by authors.

Note: The left panel shows NIR camera footage with person detection (bounding box) and skeletal pose overlay extracted in real-time. The right panel illustrates the privacy enforcement mechanism: the original video frame is immediately deleted after processing, while only body keypoint coordinates (17 COCO format landmarks) are retained for fall detection analysis. This architectural constraint renders facial recognition structurally impossible by eliminating facial data at capture time rather than relying on policy-based protections.

Figure 1. Privacy-by-Design Mechanism: Immediate Frame Disposal with Pose-Only Storage

1.2 Hardware Design

The proposed system targets middle-income Cambodian households through the following component selection:

- **Cameras:** 4× RGB cameras with 850nm IR night vision (Hikvision DS-2CD1343G2-IUF, \$63 per camera = \$252 total). Four cameras enable 360-degree coverage through 90-degree spacing.
- **Edge processor:** NVIDIA Jetson Orin Nano 8GB (\$250). Sufficient GPU capability for real-time pose estimation with 8GB RAM is adequate for multi-camera processing.
- **Accessories:** Storage (microSD card), uninterruptible power supply (UPS), ethernet cables, mounting hardware (\$170 total).

Total system cost: \$672, one-time investment, zero recurring fees.

1.3 Software Pipeline

The software architecture employs a two-stage processing pipeline:

Person detection via YOLOv8n (Nano variant, chosen for computational efficiency). YOLO identifies person bounding boxes within each frame, providing region-of-interest (ROI) coordinates for focused processing in subsequent pose estimation.

Pose estimation employs MediaPipe, extracting 33 skeletal landmarks using a two-stage detector-tracker architecture. This system uses only the 17-keypoint COCO body subset, which excludes face landmarks by design. The COCO subset includes shoulders, elbows, wrists, hips, knees, ankles, and torso keypoints, omitting the facial, hand, and foot landmarks available in MediaPipe's full 33-landmark model.

Two **processing modes** underwent validation testing: (1) Baseline configuration applies MediaPipe pose estimation directly to full frames without preliminary person detection; (2) Integrated configuration chains YOLOv8n person detection before MediaPipe pose estimation, adding computational overhead but improving detection accuracy. Both configurations exceed real-time requirements (15+ FPS).

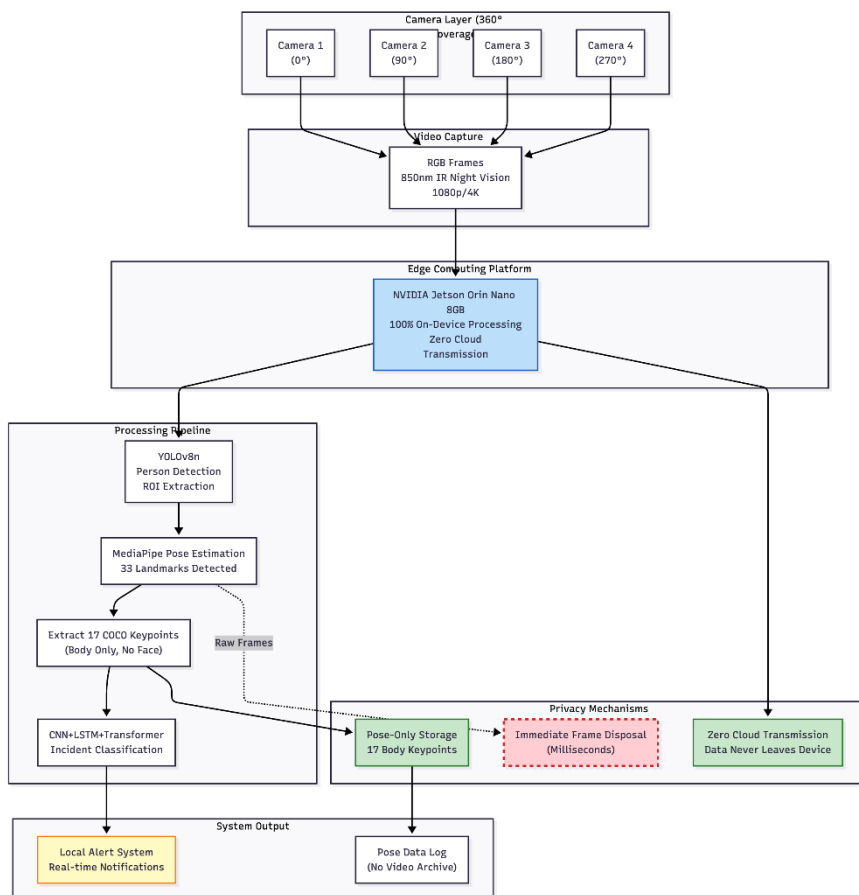
2. NIR Camera Compatibility Validation

This validation tests whether MediaPipe pose estimation, trained predominantly on visible-spectrum RGB imagery, maintains adequate performance on 850nm near-infrared footage from affordable security cameras.

2.1 Validation Dataset

MediaPipe pose estimation was validated on 20 commercial 850nm NIR security camera videos assembled from publicly available CCTV demonstration footage. The 20 videos provide:

- **Source diversity:** Professional CCTV demo footage from multiple manufacturers including Hikvision, EZviz, and other commercial brands marketing 850nm IR night vision.
- **Environmental diversity:** Indoor and outdoor environments, lighting conditions ranging from complete darkness (IR-only) to mixed visible-IR scenarios, different camera types (dome, turret, bullet).
- **Resolution range:** Both 1080p and 4K footage, reflecting affordable IP camera capabilities.



Note: Architecture enforces privacy through two mechanisms: (1) 100% on-device edge processing with zero cloud transmission, and (2) pose-only storage with immediate frame disposal (17 body keypoints, no facial landmarks, video frames deleted immediately after extraction). All computation occurs on a local edge device, eliminating network exposure and cloud data sovereignty concerns.

Figure 2. Privacy Governance-Driven System Architecture: Edge-based Elderly Safety Monitoring

- **NIR wavelength targeting:** All videos employ 850nm infrared illumination, matching the Hikvision DS-2CD1343G2-IUF camera specification.

- **Dataset limitations:** 20 videos constitute preliminary validation, not a statistically powered evaluation. Videos show general human activity rather than elderly-specific movement patterns or actual fall incidents. No ground truth pose annotations exist, requiring detection rate metrics and confidence scores rather than keypoint localization accuracy measurements.

2.2 Validation Procedure

Two processing pipelines underwent testing on the same 20 NIR videos:

Baseline pipeline: MediaPipe pose estimation applied directly to full video frames.

Integrated pipeline: YOLOv8n person detection followed by ROI extraction and cropping, with MediaPipe pose estimation applied to the cropped person region.

For each video and pipeline configuration, validation collected: (1) Keypoint detection rate (percentage of 33 landmarks successfully detected); (2) Confidence scores (average across all detected keypoints); (3) False negative rate (percentage of frames containing persons where pose detection failed entirely); (4) Processing speed (frames per second, measured on standard hardware).

2.3 Validation Scope

This validation tests: Can MediaPipe detect poses on 850nm NIR footage from affordable security cameras? It does not validate fall detection accuracy, incident classification performance, or system reliability in actual deployment conditions. The distinction between “pose detection capability” and “fall detection accuracy” proves critical for interpreting this study’s contributions.

3. Cost-Effectiveness Analysis

3.1 Comparative Cost Model

The edge-based system’s cost was compared against the Kami Fall Detect Camera (Upadhyay 2025), a commercially available cloud-based elderly fall detection product, over a 3-year period:

Edge-based system: Hardware cost \$672 one-time, recurring fees \$0, 3-year total \$672.

Cloud-based alternative: Hardware costs \$99, mandatory subscription \$45/month × 36 months = \$1,620, 3-year total \$1,719.

Cost difference: \$1,047 savings over 3 years, representing 61% cost reduction.

Breakeven point: Month 13 (Month 1 of Year 2). The edge system's higher upfront cost is offset by zero recurring fees, with cumulative costs crossing at 12.7 months.

3.2 Market Accessibility Analysis

Cambodia Socio-Economic Survey (CSES) household income data (National Institute of Statistics (Cambodia) 2021) defines market segments. The 4th and 5th quintiles (middle-income and upper-middle-income urban households) earn \$870-\$1,622 per month, representing 40% of urban households.

The \$672 system cost represents 0.41-0.77 months of income for these households, positioning the system as a significant but plausibly affordable one-time purchase. By contrast, the Kami cloud alternative's \$45/month subscription represents 2.8-5.2% of monthly income—a recurring payment that compounds over time.

Cambodia's elderly population (60+ years) will reach approximately 2.1 million by 2030 (United Nations, Department of Economic and Social Affairs, Population Division 2015). Estimating that 8-12% fall within middle-income urban households with financial capacity and motivation to adopt fall detection technology yields 168,000-252,000 potential users.

4. Evaluation Metrics

NIR camera compatibility metrics: Keypoint detection rate (percentage of 33 MediaPipe landmarks detected per frame), average confidence (mean confidence score on a 0-1 scale), false negative rate (percentage of frames with persons where pose detection failed), and processing speed (FPS, determining real-time performance).

Cost-effectiveness metrics: 3-year total cost, breakeven point (month when cumulative costs equal), cost reduction percentage, market reach (estimated number and percentage of elderly population within affordability threshold).

5. Study Scope and Limitations

5.1 What we validate:

- **NIR camera compatibility:** MediaPipe pose estimation works on 850nm infrared footage, addressing the primary technical uncertainty for privacy-preserving monitoring using affordable security cameras.
- **Cost-effectiveness comparison:** Edge architecture (\$672) costs 61% less than cloud alternatives (\$1,719) over 3 years.
- **Privacy architecture:** Storing only 17 body keypoints with immediate frame disposal makes facial identification structurally impossible.

5.2 What we do NOT validate:

- **Fall detection accuracy:** This study does not test whether the system accurately detects falls or achieves acceptable false positive/negative rates. Fall detection requires validation on benchmark datasets with ground truth fall annotations.
- **Real-world deployment:** Testing on commercial CCTV demo videos differs from continuous monitoring in actual elderly homes. Movement patterns and fall dynamics of elderly individuals differ from those of the general adult population in camera demos.
- **Hardware-specific performance:** Processing speed measurements come from standard hardware testing, not actual Jetson Orin Nano deployment.
- **User acceptance:** Elderly user and family acceptance of continuous pose monitoring remains empirically untested.
- **Coverage adequacy:** Whether the 4-camera configuration provides sufficient coverage for typical Cambodian middle-income homes remains unvalidated.
- **System reliability:** Long-duration performance over weeks and months of continuous operation exceeds this preliminary study's timeline.

This design study demonstrates how privacy governance principles inform architectural decisions, with technical validation limited to the component most uncertain for the proposed approach: NIR camera compatibility.

IV. Results

1. NIR Camera Compatibility

MediaPipe pose estimation was validated on 20 commercial 850nm NIR security camera videos representing diverse manufacturers, environments, and camera configurations.

1.1 Integrated Pipeline Performance

The integrated YOLO+MediaPipe pipeline demonstrated the following performance across the 20-video validation set:

- **Keypoint detection rate:** 91.3% (30.1 of 33 MediaPipe landmarks detected on average per frame). Missing landmarks concentrated in hands and feet (extremities), while core body keypoints (shoulders, hips, torso) achieved near-perfect detection. For fall detection purposes, core body keypoint reliability matters more than extremity precision.

- **Average confidence:** 0.868 on MediaPipe’s 0-1 confidence scale. Confidence scores above 0.85 generally indicate reliable keypoint localization. Confidence scores ranged from 0.736 to 0.905 across the 20 videos, with all videos meeting the 0.70 threshold.
- **False negative rate:** 12.3% of frames containing visible persons resulted in complete pose detection failure (zero keypoints detected). This represents approximately 2.5 failed frames per second at 20 FPS processing. This rate reflects pose detection failure on NIR footage, frames where the pipeline extracted no keypoints whatsoever, not fall event classification accuracy; the distinction is critical for interpreting this study’s scope, which does not validate whether the system detects falls.
- **Processing speed:** 20.53 FPS averaged across the 20 videos, measured on standard hardware. This exceeds the 15 FPS threshold conventionally considered “real-time” for continuous monitoring applications.

Table 1. NIR Camera Compatibility: Integrated Pipeline Performance on 20 Commercial 850nm NIR Videos

Metric	Mean	Range	Target	Status
Keypoint detection rate (%)	91.3	73.8-98.9	≥90	✓
Average confidence (0-1 scale)	0.868	0.736-0.905	≥0.70	✓
False negative rate (%)	12.3	0.0-70.2	<15	✓
Processing speed (FPS)	20.53	-	≥15	✓
Pose coverage (%) ^a	86.0	-	-	-

Note: ^aPose coverage = percentage of frames with at least one keypoint detected. Validation conducted on 20 commercial CCTV demonstration videos (850nm NIR wavelength) from multiple manufacturers including Hikvision, EZviz, and other brands. Videos included diverse camera types (dome, turret, bullet, wide-angle) and environments (indoor corridors, outdoor yards, retail spaces, residential settings). All metrics met acceptance criteria, confirming MediaPipe pose estimation compatibility with affordable 850nm NIR security cameras.

1.2 Pipeline Comparison

The baseline MediaPipe-only pipeline (pose estimation applied directly to full frames without preliminary person detection) provides a comparative reference:

Baseline pipeline performance: Keypoint detection rate 85.6% (28.3 of 33 landmarks), average confidence 0.833, false negative rate 20.5%, processing speed 47.37 FPS.

Performance differentials (integrated vs. baseline):

- Detection accuracy: 91.3% vs. 85.6% = 5.7 percentage point improvement (6.7% relative improvement)
- Pose coverage: 86.0% (frames with at least one keypoint detected) vs. 63.8% baseline, representing 22.2 percentage point improvement in coverage reliability
- Confidence: 0.868 vs. 0.833 = 0.035 point improvement
- Speed cost: 20.53 FPS vs. 47.37 FPS = 2.3× slower processing for integrated pipeline

The integrated YOLO+MediaPipe approach was selected for the proposed system architecture despite 2.3× slower processing speed. The 5.7 percentage point detection accuracy improvement and 22.2 percentage point better pose coverage provide meaningful reliability gains for fall detection, where false negatives (missing actual falls) carry potentially fatal consequences. Both pipelines process faster than a 15 FPS real-time threshold, so the speed differential represents a performance margin rather than a functionality threshold.

Table 2. Processing Pipeline Comparison: Baseline vs. Integrated Architecture

Metric	Baseline (MediaPipe Only)	Integrated (YOLO+MediaPipe)	Change
Processing speed (FPS)	47.37	20.53	-56.7%
Keypoint detection rate (%)	85.6	91.3	+5.7 pp ^a
Average confidence	0.833	0.868	+0.035
Pose coverage (%)	63.8	86.0	+22.2 pp
False negative rate (%)	20.5	12.3	-8.2 pp
Speed-Accuracy Ratio ^b	2.31	1.00	-

^app = percentage points (absolute difference, not relative change).

^bSpeed-Accuracy Ratio = (FPS / FPS_"integrated") /

(Detection/Detection_"integrated"). Values >1 indicate faster processing; the integrated pipeline is normalized to 1.0.

Note: Both pipelines were tested on identical 20 NIR videos. Baseline applies MediaPipe pose estimation directly to full frames. Integrated uses YOLOv8n person detection followed by ROI-cropped MediaPipe processing. Integrated

pipeline selected for proposed system despite 2.3× slower processing speed due to 5.7 percentage point higher detection accuracy and 22.2 percentage point better pose coverage, prioritizing reliability over efficiency for safety-critical fall detection applications.

1.3 Performance Variation

Performance varied across the 20 individual videos:

- **Keypoint detection range:** 73.8% to 98.9%. The highest-performing video (98.9% detection) showed stable indoor footage from a dome camera, while the lowest-performing video (73.8% detection) featured outdoor scenes with motion blur and occlusion. Elderly monitoring in indoor home environments (the target use case) corresponds to scenarios achieving higher performance.
- **Environmental factors:** The validation dataset included diverse environments (indoor corridors, outdoor yards, retail spaces, residential settings), with indoor scenarios under stable lighting achieving the highest detection rates.
- **Camera type diversity:** The 20-video validation tested multiple camera types including dome, turret, bullet, wide-angle (180°), and hidden spy cameras across both 1080p and 4K resolutions, confirming 850nm NIR compatibility across diverse commercial CCTV equipment.

These variations suggest that NIR camera compatibility depends on environmental factors, camera placement, and specific camera models. Practitioners implementing similar systems should conduct camera-specific validation.

2. Cost-Effectiveness Analysis Results

2.1 Total Cost Comparison

The edge-based system demonstrates substantial cost advantages over cloud-based fall detection alternatives across a 3-year deployment period:

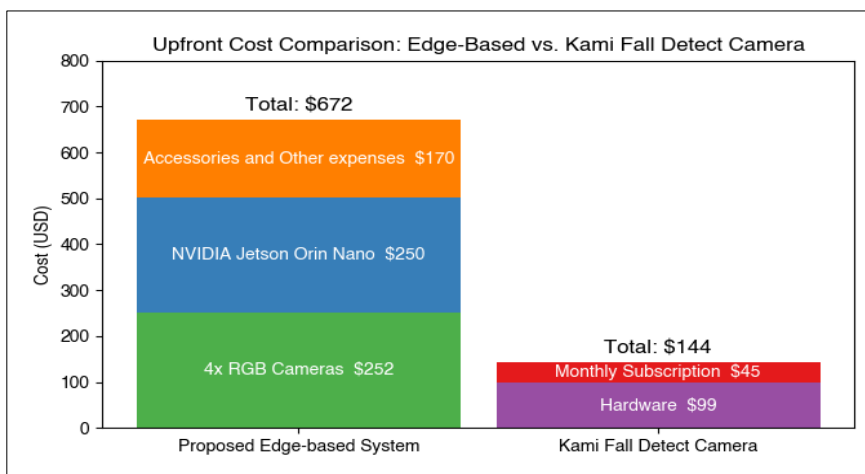
- Edge-based system (proposed architecture): Hardware components (\$252 cameras + \$250 Jetson Orin Nano + \$170 accessories) = \$672 initial investment, recurring monthly fees \$0, 3-year total cost \$672.
- Cloud-based alternative (Kami Fall Detect Camera): Hardware \$99, recurring subscription \$45/month × 36 months = \$1,620, 3-year total cost \$1,719.
- Cost differential: \$1,047 savings over 3 years, representing 61% cost reduction. The edge system costs less than 40% of the cloud alternative over the 3-year period.

- Breakeven analysis: The edge system's higher upfront cost (\$672 vs. \$99) reverses over time as avoided subscription fees accumulate. Breakeven occurs at Month 13 (early in Year 2), after which savings increase by \$45 per month. For elderly monitoring where continuous deployment over multiple years seems likely, the long-term cost advantage proves more relevant than the initial price differential.

2.2 Market Accessibility Implications

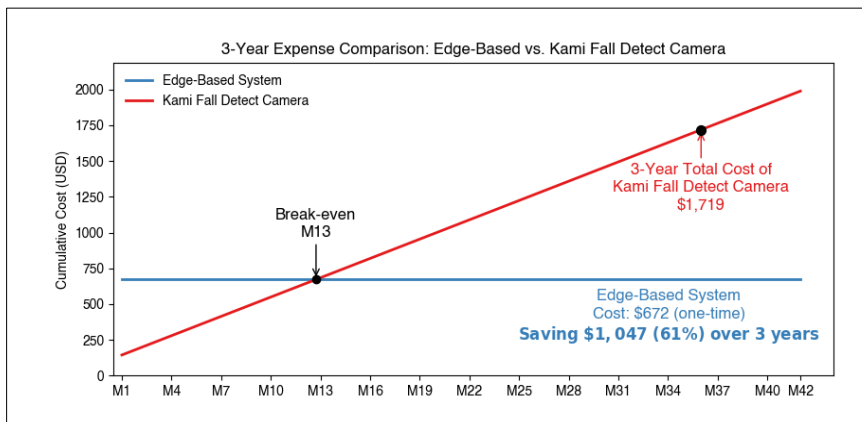
The \$672 edge system represents 0.41-0.77 months of household income for middle-income Cambodian urban households (4th-5th quintiles earning \$870-\$1,622/month), positioning the technology as a significant but potentially affordable one-time purchase. By contrast, the \$1,719 cloud system total cost represents 1.06-1.98 months of income—nearly double the edge system's income burden.

The subscription structure creates an ongoing payment burden. The \$45/month Kami subscription represents 2.8-5.2% of monthly household income, every month, indefinitely. Recurring payments pose adoption barriers beyond their absolute cost, particularly when payments must continue indefinitely for service access.



Note. Edge-based system requires higher initial investment (\$672: \$252 cameras + \$250 edge processor + \$170 accessories) but zero recurring fees. Cloud-based Kami Fall Detect Camera shows a lower upfront cost (\$144: \$99 hardware + \$45 first month subscription) but requires continuous monthly payments for fall detection functionality. Upfront cost difference (\$528) reverses by Month 13 when cumulative subscription fees exceed the edge system's one-time investment.

Figure 3. Upfront Cost Comparison: Edge-Based vs. Cloud-Based Elderly Safety Monitoring Systems



Note: Cumulative cost analysis over a 36-month deployment period. Edge-based system (blue line) maintains a constant \$672 cost after initial investment. Cloud-based system (red line) accumulates subscription fees linearly at \$45/month. Breakeven occurs at Month 13 when cumulative costs equalize. By Year 3 end, edge architecture achieves \$1,047 savings (61% cost reduction), with the advantage continuing to compound in subsequent years. For elderly monitoring applications requiring multi-year continuous operation, the economic case for edge architecture strengthens over time.

Figure 4. Three-Year Cumulative Cost Trajectory: Edge Architecture vs. Cloud Subscription Model

Table 3. Three-Year Cost Comparison: Edge-Based vs. Cloud-Based Elderly Safety Monitoring Systems

System Component	Year 1	Year 2	Year 3	Total
Edge-Based System (Proposed)				
4 × RGB cameras (850nm IR)	\$252	\$0	\$0	\$252
NVIDIA Jetson Orin Nano 8GB	\$250	\$0	\$0	\$250
Accessories (storage, UPS, cables)	\$170	\$0	\$0	\$170
Subtotal	\$672	\$0	\$0	\$672
Cloud-Based System (Kami Fall Detect Camera)				
Hardware	\$99	\$0	\$0	\$99
Cloud subscription (\$45/month)	\$540	\$540	\$540	\$1,620
Subtotal	\$639	\$540	\$540	\$1,719
Cost Savings (Edge vs. Cloud)	-\$33	+\$540	+\$540	+\$1,047
Cost Reduction Percentage	-5%	100%	100%	61%

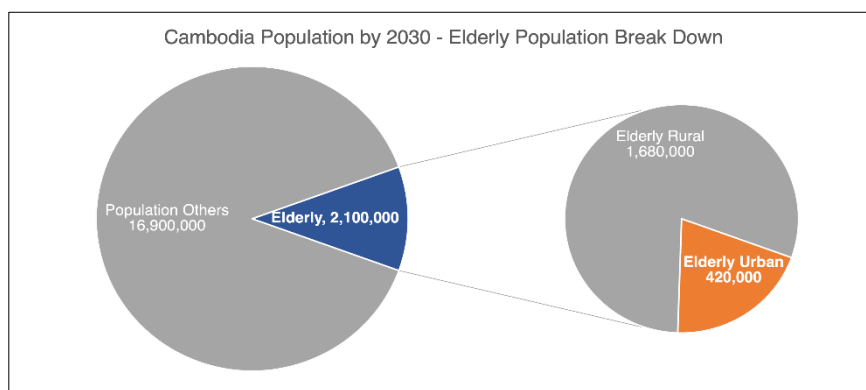
Note: Edge-based system has a higher upfront cost but zero recurring fees, achieving breakeven at Month 13 (early Year 2). The cloud system requires a continuous

subscription for fall detection functionality. Cost savings calculated as: $(\$1,719 - \$672) / \$1,719 = 61\%$ reduction over 3 years.

Market reach estimation: Applying affordability thresholds to Cambodia’s projected 2030 elderly population:

- Total elderly population (60+ years): 2.1 million
- Urban elderly (20% urbanization rate for elderly): 420,000
- Middle-income urban elderly (40% of urban households): 168,000 potential adopters
- Conservative estimate (60% adoption factor): 252,000 potential adopters

The 168,000-252,000 range represents 8-12% of Cambodia’s projected elderly population. By comparison, the cloud system’s higher cost might contract market reach to approximately the 5th quintile only: 126,000 potential adopters (6% of the elderly population), representing a 25-50% smaller market than the edge system reach.



Source: Authors, based on UN DESA (2015) and Cambodia Socio-Economic Survey 2021
 Note: The left pie chart shows Cambodia’s projected total population for 2030 (19 million), with the elderly segment (2.1 million, 11%) highlighted. The right pie-of-pie chart breaks down the elderly population into urban (420,000, 20%) and rural (1,680,000, 80%) segments. The target market comprises middle-income urban elderly households (4th-5th quintiles), representing 168,000-252,000 potential adopters (8-12% of the elderly population). Rural elderly and low-income urban populations require alternative deployment approaches beyond market-based technology adoption.

Figure 5. Cambodia Elderly Population Segmentation by 2030: Urban-Rural Distribution and Target Market

Table 4. Market Accessibility: Estimated Reach in Cambodia’s Elderly Population (2030 Projection)

Population Segment	Count	% of Elderly
Total elderly population (60+ years, 2030)	2,100,000	100%
Urban elderly (20% urbanization) ^a	420,000	20%
4th–5th quintile (middle-income) ^b	168,000	8%
Including 3rd quintile (family pooling) ^c	252,000	12%
Affordability Context		
4th quintile monthly income	\$870/month	
5th quintile monthly income	\$1,622/month	
Edge system cost (% of monthly income)	0.41-0.77 months	
Cloud system monthly cost (% of income)	2.8-5.2% monthly	

Source: Authors, based on Elderly population: UN DESA (2015). Income: Cambodia Socio-Economic Survey 2021. System costs: Validation analysis. Urbanization rate (20%) is stated assumption.

Note:

^aElderly urbanization (20%) is lower than the general population (40%) due to rural concentration of older adults.

^bEach quintile represents 20% of households. 4th + 5th quintile = 40% of urban elderly: $420,000 \times 40\% = 168,000$.

^cAdding 3rd quintile (20%) to 4th-5th quintiles (40%) = 60% of urban elderly: $420,000 \times 60\% = 252,000$. The third quintile (\$648/month) faces a higher burden but may access through family pooling.

Edge system (\$672) expands market reach to 168,000-252,000 elderly (8-12% of total).

Cloud alternative (\$1,719) likely limited to 5th quintile only (84,000 elderly, 4% of total) due to higher cost burden.

Table 5. Cambodian Household Income Quintiles and System Affordability Assessment (Urban Households, 2021)

Income Quintile	Monthly Income	Affordability	Market Segment
Top 20% (5th quintile)	\$1,622/month	✓ Highly feasible	168,000 elderly (8% of total)
Fourth 20% (4th quintile)	\$870/month	✓ Feasible with savings	Included in the target market
Middle 20% (3 rd quintile)	\$648/month	! Requires pooling	Beyond current market reach
Lower 40% (1 st -2 nd quintile)	<\$500/month	x Cannot afford unless subsidized	Requires subsidies or policy support

Source: Cambodia Socio-Economic Survey (CSES) 2021 (National Institute of Statistics (Cambodia) 2021).

Note:

^aFamily pooling: Multiple household members contributing to purchase cost, common practice in Cambodian extended families.

Edge system at \$672 represents 0.41 months' income for the 5th quintile and 0.77 months for the 4th quintile, positioning as a significant but affordable one-time purchase for middle-income households. Third quintile (\$648/month) faces a higher burden (1.04 months' income) but may access through family pooling. Lower quintiles require policy interventions beyond market mechanisms for technology access.

2.3 Cost Model Limitations

The analysis omits electricity consumption, internet connectivity costs, potential maintenance or replacement costs, and installation labor. Including these factors would modify absolute costs but likely preserve relative cost advantage for edge system. Some cloud services offer annual subscription discounts. The 3-year analysis period assumes households maintain the system continuously for 36 months; if systems remain functional beyond 3 years, the edge system's cost advantage compounds further.

The reported cost comparison (\$672 vs. \$1,719, 61% reduction) should be interpreted as indicative of direction rather than precise magnitude. The finding that edge architecture provides substantial cost advantages over subscription-based cloud monitoring remains valid, though exact savings depend on deployment specifics and household usage patterns.

V. Discussion

1. Governance Implications

1.1 Privacy Governance Through Edge Architecture

This study demonstrates how privacy governance principles can drive architectural design decisions rather than being retrofitted as compliance measures. Edge-based processing eliminates cloud transmission requirements through system constraints rather than operational promises, a difference that becomes critical when considering enforcement challenges that plague privacy-by-policy approaches.

The architecture achieves facial anonymity by design: pose-only storage (17 body keypoints in COCO format) renders facial recognition impossible, not merely prohibited. Systems that capture facial data and promise not to use it for identification require trusting that policies will be followed. By contrast, storing only skeletal coordinates makes facial identification structurally impossible because facial data never existed in stored form. Raw video frames undergo

immediate deletion after keypoint extraction (milliseconds), eliminating retrospective re-identification risks.

The validation results (91.3% keypoint detection on 850nm NIR footage) confirm technical feasibility without performance compromise. The edge-pose-only approach achieves privacy through architecture while maintaining pose estimation performance comparable to visible-spectrum applications. In contexts where data protection regulations provide limited enforcement, such as Cambodia, privacy protection must emerge from system design. Edge architecture with pose-only storage demonstrates technical pathway for governance-driven design that doesn't rely on external enforcement—the architecture is the governance.

One residual privacy risk warrants acknowledgment. Skeletal coordinate sequences can encode gait patterns sufficient for biometric identification even without facial data (Teepe et al. 2021), and appearance-level anonymization alone does not neutralize this vulnerability (Weiß et al. 2026). For the current architecture, this means that pose-only storage eliminates facial identity but preserves the movement signatures that gait biometrics exploit. The architecture's core privacy contribution—eliminating video retention and cloud transmission—remains intact, but skeletal data is not biometrically inert. Future iterations should evaluate whether gait obfuscation or differential privacy techniques can mitigate this residual risk.

Local device security represents a related gap. The edge architecture eliminates cloud transmission risks—data that never leaves the local device cannot be intercepted in transit—but security of the Jetson platform itself was outside this study's design scope. Physical access controls, storage encryption for retained pose logs, and firmware integrity protections were neither designed nor evaluated here. These concerns are distinct from cloud privacy risks and warrant attention in future deployment work.

1.2 Accessibility Governance Through Cost Optimization

Edge architecture yields economic co-benefits beyond privacy protection. Eliminating cloud infrastructure reduces 3-year total cost by 61% (\$672 vs. \$1,719), expanding market accessibility in ways that matter for health equity.

The zero-subscription model removes recurring payment barriers particularly challenging for middle-income Cambodian households earning \$870-\$1,622/month, where \$45/month subscription represents 2.8-5.2% of monthly income indefinitely. The estimated market reach of 168,000-252,000 elderly individuals (8-12% of Cambodia's projected 2030 elderly population) represents meaningful expansion beyond high-income early adopters who can afford subscription-based cloud monitoring.

Edge-based architecture expands accessibility beyond cloud alternatives, though substantial portions of elderly population remain unreachable. Low-

income and rural populations require alternative deployment models. Technology design alone cannot solve health equity challenges rooted in economic inequality; it can avoid exacerbating inequality through pricing structures that systematically exclude middle- and low-income populations. The zero-subscription edge model represents harm reduction—eliminating recurring payment barriers—rather than universal accessibility.

1.3 Context-Specific Design for Developing Countries

The system design incorporates regional epidemiological evidence: three incident types (falls while standing/walking, falls from bed/chair, abnormal sit-to-stand transitions) were prioritized based on Thai elderly fall data (Maiyapakdee et al. 2025) and Chinese long-term care studies (Li and Shi 2022). This evidence-based incident prioritization contrasts with technology-driven approaches that detect whatever algorithms can most easily classify.

Hardware selection (850nm NIR cameras at \$63 each, edge processor) addresses Cambodia's middle-income market constraints while maintaining technical performance. The cameras represent approximately one-tenth the cost of research-grade RGB-D sensors frequently employed in fall detection studies, fundamentally altering market accessibility. This demonstrates a governance-driven approach for resource-constrained contexts where off-the-shelf solutions developed for high-income markets prove economically inaccessible.

The system targets indoor home environments exclusively. This design boundary reflects the deployment context: strong family caregiving norms across Southeast Asian societies favor home-based elderly care over institutional alternatives (Romli et al. 2017), and the incident types prioritized—falls while standing or walking, falls from bed or chair, abnormal sit-to-stand transitions—are predominantly indoor events (Maiyapakdee et al. 2025). Continuous 24/7 monitoring also requires fixed infrastructure, most practically deployed within the home. Extending coverage to outdoor environments would introduce fundamentally different technical constraints (variable lighting, wider coverage areas, weatherproofing) and cost structures, constituting a separate design problem.

This study's approach is what innovation theory terms technology-push: beginning from governance principles and technical constraints to produce an architecture, then evaluating its feasibility and cost-effectiveness (Mowery and Rosenberg 1979). The complementary perspective, demand-pull, asks whether intended users actually want and would adopt the system. That question remains unanswered here. Whether elderly Cambodians and their families perceive a need for continuous pose-based monitoring, whether they trust AI surveillance in the home, and how family caregiving dynamics shape adoption decisions are demand-side factors that this design study does not investigate. The distinction matters: technical feasibility and affordability are necessary but not sufficient

conditions for adoption, and future work should engage directly with demand-pull dynamics through empirical user acceptance research.

2. Technical Validation Contributions

2.1 NIR Camera Compatibility

The validation of MediaPipe pose estimation on 850nm NIR footage (91.3% keypoint detection across 20 commercial CCTV videos) fills a literature gap. The specific question of whether pose estimation algorithms trained on visible-spectrum data maintain performance on 850nm infrared imagery from affordable CCTV cameras sits at the intersection of computer vision, healthcare AI, and assistive technology research domains. This study provides empirical evidence that pose estimation works on NIR footage, with performance characteristics enabling 24/7 elderly monitoring without facial recognition technology.

The practical contribution is deployment guidance: affordable NIR cameras prove viable for pose-based applications, though practitioners should validate camera-specific performance given the 73.8-98.9% detection range observed. Environmental factors matter (stable indoor lighting achieved the highest performance, while outdoor scenes with motion blur showed the lowest), camera placement affects results, and manufacturer differences exist.

2.2 Pipeline Trade-off Analysis

The integrated YOLO+MediaPipe approach achieved 5.7 percentage point higher keypoint detection and 22.2 percentage point better pose coverage, but processed 2.3× slower than baseline (20.53 FPS vs. 47.37 FPS). The decision to select the integrated pipeline despite the speed penalty reflects the governance principle: when safety-critical applications face accuracy-speed tradeoffs and both configurations exceed minimum requirements (both >15 FPS), prioritizing reliability over efficiency embodies different values than commercial applications optimizing for computational cost reduction.

3. Limitations

Validation scope: NIR camera compatibility (pose estimation on 850nm infrared footage) and cost-effectiveness (edge vs. cloud expense comparison) were validated. Fall detection accuracy, incident classification performance, real-world deployment reliability, user acceptance, coverage adequacy, and long-term system maintenance were not validated. The gap between “pose estimation works on NIR cameras” and “the system accurately detects falls in elderly homes” spans multiple validation dimensions this preliminary study does not address.

Testing environment: Commercial CCTV demonstration footage shows general adult populations rather than elderly individuals with age-specific movement patterns. Demo videos feature normal daily activities rather than fall incidents. Whether 91.3% keypoint detection on general populations generalizes to elderly subjects and fall events remains empirically unverified. A further limitation concerns mobility aids: older adults who use walkers, rollators, or canes present an additional challenge, as assistive devices can occlude lower-body keypoints and pose estimation applied to elderly populations with such aids shows measurably degraded performance (Ng et al. 2020). Validation on elderly-specific datasets, including subjects using assistive devices, is necessary before deployment claims can be made with confidence.

Hardware deployment: Processing speeds (20.53 FPS integrated, 47.37 FPS baseline) were measured on standard hardware, not on the target NVIDIA Jetson Orin Nano 8GB. Whether Jetson achieves similar speeds for 4-camera processing requires hardware-specific validation.

Market accessibility: The 168,000-252,000 user estimate represents economic feasibility (households with financial capacity) but not actual adoption. Perceived need, technology trust, family dynamics, and alternative options beyond cost determine technology adoption.

Indoor scope: The system addresses indoor home environments only. Outdoor falls, which also pose risks to elderly populations, are outside this study's scope and would require a substantially different architectural approach.

Generalizability: Cambodia-specific characteristics (family caregiving traditions, middle-income household income ranges, limited data protection regulation) shape problem definition and solution constraints. The governance-driven design approach might transfer across contexts, but specific architectural choices require context-specific adaptation.

4. Future Directions

Benchmark dataset validation. Fall detection accuracy requires validation on standard datasets (MCF, LE2I, UP-Fall) with ground truth annotations to establish whether pose-based incident detection achieves acceptable accuracy for deployment.

Custom dataset collection. Collecting a Southeast Asian elderly dataset (180 videos featuring age-appropriate actors across three lighting conditions) would address Western dataset bias while enabling lighting-condition-specific algorithm optimization.

Hardware validation. Deploying the complete pipeline on NVIDIA Jetson Orin Nano and measuring actual 4-camera processing speeds would validate

deployment feasibility claims. Sustained-operation testing (hours to days) would identify thermal throttling or reliability issues.

Expanded detection capabilities. Beyond the three current incident types, elderly safety encompasses prolonged floor lying, bathroom incidents, wandering behavior (dementia patients), and abnormal inactivity patterns.

Policy framework development. Broader policy questions remain: informed consent processes for elderly monitoring, data control and alert settings, family responsibilities, and healthcare provider integration of home monitoring data.

Longitudinal adoption studies. Understanding actual adoption patterns, sustained usage rates, user acceptance evolution, and long-term value proposition requires studies following households over months to years.

VI. Conclusion

1. Key Findings

This study demonstrates privacy governance-driven architectural design through a Cambodia elderly monitoring design study, showing that starting with privacy principles and cost constraints rather than technical optimization yields different system architectures with implications for governance implementation and market accessibility.

Technical feasibility. Edge-based pose estimation achieves 91.3% keypoint detection on 850nm NIR footage from affordable security cameras, validating 24/7 privacy-preserving monitoring capability without expensive visible-light or depth sensors. This establishes deployment feasibility: affordable cameras (\$63 each) can support pose-based fall detection, obviating the need for cost-prohibitive sensor alternatives that exclude middle-income markets. The finding comes with caveats (12.3% false negatives, 73.8-98.9% range across individual videos, testing on general populations rather than elderly subjects), but NIR compatibility works adequately for this application.

Cost-effectiveness. Edge architecture reduces 3-year total cost by 61% compared to cloud-based fall detection alternatives (\$672 vs. \$1,719), expanding market reach to an estimated 168,000-252,000 elderly individuals in middle-income Cambodian households. The cost comparison reflects the structural difference between one-time hardware investment and perpetual subscription fees, with the edge system achieving breakeven at Month 13. The zero-subscription model eliminates recurring payment barriers that disproportionately affect households with variable income streams.

Design trade-offs. The selection of integrated YOLO+MediaPipe pipeline over simpler baseline approach—despite 2.3× slower processing speed—

reflects safety-critical priority where accuracy trumps efficiency when both configurations exceed minimum requirements. The integrated pipeline's 5.7 percentage point higher keypoint detection and 22.2 percentage point better pose coverage provide reliability improvements for applications where false negatives (missed fall detections) carry potentially fatal consequences.

Governance-driven design. Privacy governance principles informed architectural decisions from inception (edge processing, pose-only storage, immediate frame disposal) rather than being retrofitted post-deployment, demonstrating Privacy by Design principles in the healthcare AI context. The architecture enforces privacy through system constraints: facial recognition becomes impossible when no facial data exists, data breaches cannot expose video archives that were never created, and cloud transmission risks vanish when processing never leaves local devices. This differs from privacy-by-policy approaches that capture sensitive data and promise not to misuse it.

2. Implications for Practice

2.1 For Healthcare AI Developers

NIR camera compatibility requires empirical validation before deployment: the 73.8-98.9% detection range demonstrates that performance varies substantially by camera model, environmental conditions, and deployment specifics. Developers should conduct camera-specific validation sampling of anticipated deployment conditions.

Integrated pipelines achieve real-time performance on standard hardware (20.53 FPS on consumer-grade GPU), suggesting edge deployment feasibility though Jetson Orin Nano validation remains necessary. Pose-based fall detection doesn't require specialized AI workstation hardware, which matters for cost-constrained deployments.

Privacy-by-design architecture is technically feasible for elderly monitoring without sacrificing essential functionality. Pose-only storage eliminates facial recognition capability while maintaining adequate data for fall detection, demonstrating that privacy protections need not degrade application performance when embedded into architecture from inception.

2.2 For Policymakers

Privacy-by-design architecture yields economic co-benefits beyond privacy protection: the 61% cost reduction emerges from edge processing motivated by privacy governance that also eliminates cloud infrastructure expenses. Privacy and affordability need not trade off against each other.

Zero-subscription models expand healthcare technology accessibility in middle-income markets, reaching an estimated 8-12% of the elderly population. However, low-income and rural elderly require different deployment

approaches—potentially subsidized programs, community-based installations, or policy interventions—that technology design alone cannot provide.

Governance frameworks require implementation guidance specific to developing country contexts rather than assuming Western regulatory models (GDPR, HIPAA) apply universally. When data protection regulations provide limited enforcement, privacy protection must emerge from system architecture rather than compliance mechanisms.

References

- Almeida, Patricia, Carlos Santos, and Josivania Silva Farias. 2020. Artificial Intelligence Regulation: A Meta-Framework for Formulation and Governance. https://www.researchgate.net/publication/351039094_Artificial_Intelligence_Regulation_a_framework_for_governance.
- Birkstedt, Teemu, Matti Minkkinen, Anushree Tandon, and Matti Mäntymäki. 2023. "AI Governance: Themes, Knowledge Gaps and Future Agendas." *Internet Research* 33 (7): 133–67. <https://doi.org/10.1108/INTR-01-2022-0042>.
- Burns, Thomas F. 2024. "Generative AI Policy and Governance Considerations for Health Security in Southeast Asia." November 3, 2024. <https://doi.org/10.48550/arXiv.2411.14435>.
- Cavoukian, Ann, Scott Taylor, and Martin E. Abrams. 2010. "Privacy by Design: Essential for Organizational Accountability and Strong Business Practices." *Identity in the Information Society* 3 (2): 405–13. <https://doi.org/10.1007/s12394-010-0053-z>.
- Chang, Wan-Jung, Chia-Hao Hsu, and Liang-Bi Chen. 2021. "A Pose Estimation-Based Fall Detection Methodology Using Artificial Intelligence Edge Computing." *IEEE Access* 9 (September): 129965–76. <https://doi.org/10.1109/ACCESS.2021.3113824>.
- Chang-Yueh Wang, and Fang-Suey Lin. 2024. "AI-Driven Privacy in Elderly Care: Developing a Comprehensive Solution for Camera-Based Monitoring of Older Adults." *Applied Sciences* 14 (10, 10): 4150. <https://doi.org/10.3390/app14104150>.
- Chaudhuri, Shomir, Hilaire Thompson, and George Demiris. 2014. "Fall Detection Devices and Their Use with Older Adults: A Systematic Review." *Journal of Geriatric Physical Therapy* (2001) 37 (4): 178–96. <https://doi.org/10.1519/JPT.0b013e3182abe779>.
- Hatef, Elham, Sarah Scholle, Bryan Buckley, Jonathan Weiner, and John Austin. 2024. "Development of an Evidence- and Consensus-Based Digital Healthcare Equity Framework." *JAMIA Open* 7 (November): ooa136. <https://doi.org/10.1093/jamiaopen/ooae136>.
- Jalal, Ahmad, Shaharyar Kamal, and Daijin Kim. 2014. "A Depth Video Sensor-Based Life-Logging Human Activity Recognition System for Elderly Care in Smart Indoor Environments." *Sensors* 14 (7): 11735–59. <https://doi.org/10.3390/s140711735>.
- Li, Chunyan, and Chunhong Shi. 2022. "Adverse Events and Risk Management in Residential Aged Care Facilities: A Cross-Sectional Study in Hunan, China." *Risk Management and Healthcare Policy* 15 (March): 529–42. <https://doi.org/10.2147/RMHP.S351821>.
- Liu, WeiKang, Yanchun Zhang, Hong Yang, and Qinxue Meng. 2023. "A Survey on Differential Privacy for Medical Data Analysis." *Annals of Data Science*, June, 1–15. <https://doi.org/10.1007/s40745-023-00475-3>.
- Maiyapakdee, Chavinthorn, Teeraphun Kaewdok, Nontiya Homkham, Katiya Ivanovitch, and Saowanee Norkaew. 2025. "Factors Related to Home Accidents Among the Elderly in Thailand." *Social Sciences & Humanities Open* 11 (January): 101265. <https://doi.org/10.1016/j.ssaho.2024.101265>.

- Mowery, David, and Nathan Rosenberg. 1979. "The Influence of Market Demand Upon Innovation: A Critical Review of Some Recent Empirical Studies." *Research Policy* 8 (2): 102–53. [https://doi.org/10.1016/0048-7333\(79\)90019-2](https://doi.org/10.1016/0048-7333(79)90019-2).
- National Institute of Statistics (Cambodia). 2021. "Cambodia Socio-Economic Survey 2021 | GHDx." 2021. <https://ghdx.healthdata.org/record/cambodia-socio-economic-survey-2021>.
- Ng, Kimberley-Dale, Sina Mehdizadeh, Andrea Iaboni, Avril Mansfield, Alastair Flint, and Babak Taati. 2020. "Measuring Gait Variables Using Computer Vision to Assess Mobility and Fall Risk in Older Adults With Dementia." *IEEE Journal of Translational Engineering in Health and Medicine* 8: 1–9. <https://doi.org/10.1109/JTEHM.2020.2998326>.
- Richardson, Safiya, Katharine Lawrence, Antoinette M. Schoenthaler, and Devin Mann. 2022. "A Framework for Digital Health Equity." *Npj Digital Medicine* 5 (1): 119. <https://doi.org/10.1038/s41746-022-00663-0>.
- Romli, M. H., M. P. Tan, L. Mackenzie, M. Lovarini, P. Suttanon, and L. Clemson. 2017. "Falls Amongst Older People in Southeast Asia: A Scoping Review." *Public Health* 145 (April): 96–112. <https://doi.org/10.1016/j.puhe.2016.12.035>.
- Sylla, Bry, Ouedraogo Ismaila, and Gayo Diallo. 2025. "25 Years of Digital Health Toward Universal Health Coverage in Low- and Middle-Income Countries: Rapid Systematic Review." *Journal of Medical Internet Research* 27 (1): e59042. <https://doi.org/10.2196/59042>.
- Teepe, Torben, Ali Khan, Johannes Gilg, Fabian Herzog, Stefan Hörmann, and Gerhard Rigoll. 2021. "Gaitgraph: Graph Convolutional Network for Skeleton-Based Gait Recognition." In 2021 IEEE International Conference on Image Processing (ICIP), 2314–18. <https://doi.org/10.1109/ICIP42928.2021.9506717>.
- Uddin, Md. Zia, Weria Khaksar, and Jim Torresen. 2018. "Ambient Sensors for Elderly Care and Independent Living: A Survey." *Sensors (Basel, Switzerland)* 18 (7): 2027. <https://doi.org/10.3390/s18072027>.
- United Nations, Department of Economic and Social Affairs, Population Division. 2015. "Data Portal Population Division." Population Division Data Portal. 2015. <https://population.un.org/dataportal/data/indicators/70/locations/116/start/1990/end/2030/table/pivotbylocation?df=bc266e85-33a7-4e60-aaff-9a341d2efc17>.
- Upadhyay, Rishaj. 2025. "This Camera Can Detect When You Fall!" *Android Headlines*. January 9, 2025. <https://www.androidheadlines.com/2025/01/kami-visions-fall-detect-camera.html>.
- Vaiyapuri, Thavavel, E. Laxmi Lydia, Mohamed Yacin Sikkandar, Vicente García Díaz, Irina V. Pustokhina, and Denis A. Pustokhin. 2021. "Internet of Things and Deep Learning Enabled Elderly Fall Detection Model for Smart Homecare." *IEEE Access* 9 (July): 113879–88. <https://doi.org/10.1109/ACCESS.2021.3094243>.
- Wang, Xueyi, Joshua Ellul, and George Azzopardi. 2020. "Elderly Fall Detection Systems: A Literature Survey." *Frontiers in Robotics and AI* 7 (June). <https://doi.org/10.3389/frobt.2020.00071>.
- Weiß, Sarah, Christopher Bonenberger, Tobias Niedermaier, Maik Knof, and Markus Schneider. 2026. "Privacy Beyond the Face: Assessing Gait Privacy Through Realistic Anonymization in Industrial Monitoring." *Sensors* 26 (1): 187. <https://doi.org/10.3390/s26010187>.

- i.org/10.3390/s26010187.
- WHO. 2021. "Falls." April 26, 2021. <https://www.who.int/news-room/fact-sheets/detail/falls>.
- . 2025. "Ageing and Health in SEARO." October 14, 2025. <https://www.who.int/southeastasia/health-topics/ageing>.
- Williamson, Steven M., and Victor Prybutok. 2024. "Balancing Privacy and Progress: A Review of Privacy Challenges, Systemic Oversight, and Patient Perceptions in AI-Driven Healthcare." *Applied Sciences* 14 (2): 675. <https://doi.org/10.3390/app14020675>.
- Yu, Zhigang, Jiahui Liu, Mingchuan Yang, Yanmin Cheng, Jie Hu, and Xinchu Li. 2022. "An Elderly Fall Detection Method Based on Federated Learning and Extreme Learning Machine (Fed-ELM)." *IEEE Access* 10 (December): 130816–24. <https://doi.org/10.1109/ACCESS.2022.3229044>.